Compressed Air Alliance
ABN: 72 633 254 476
Email: sales@compressedairalliance.com

Data Acquisition System

Security

Compressed Air Alliance and the providers of the Data Acquisition System take the security of customers and their data very seriously. The providers of the Data Acquisition System employ a proven combination of techniques to provide enterprise-level encryption. It consists of multiple lines of defence, preventing a potential attacker who has gained control of a part of the application from accessing the whole application.

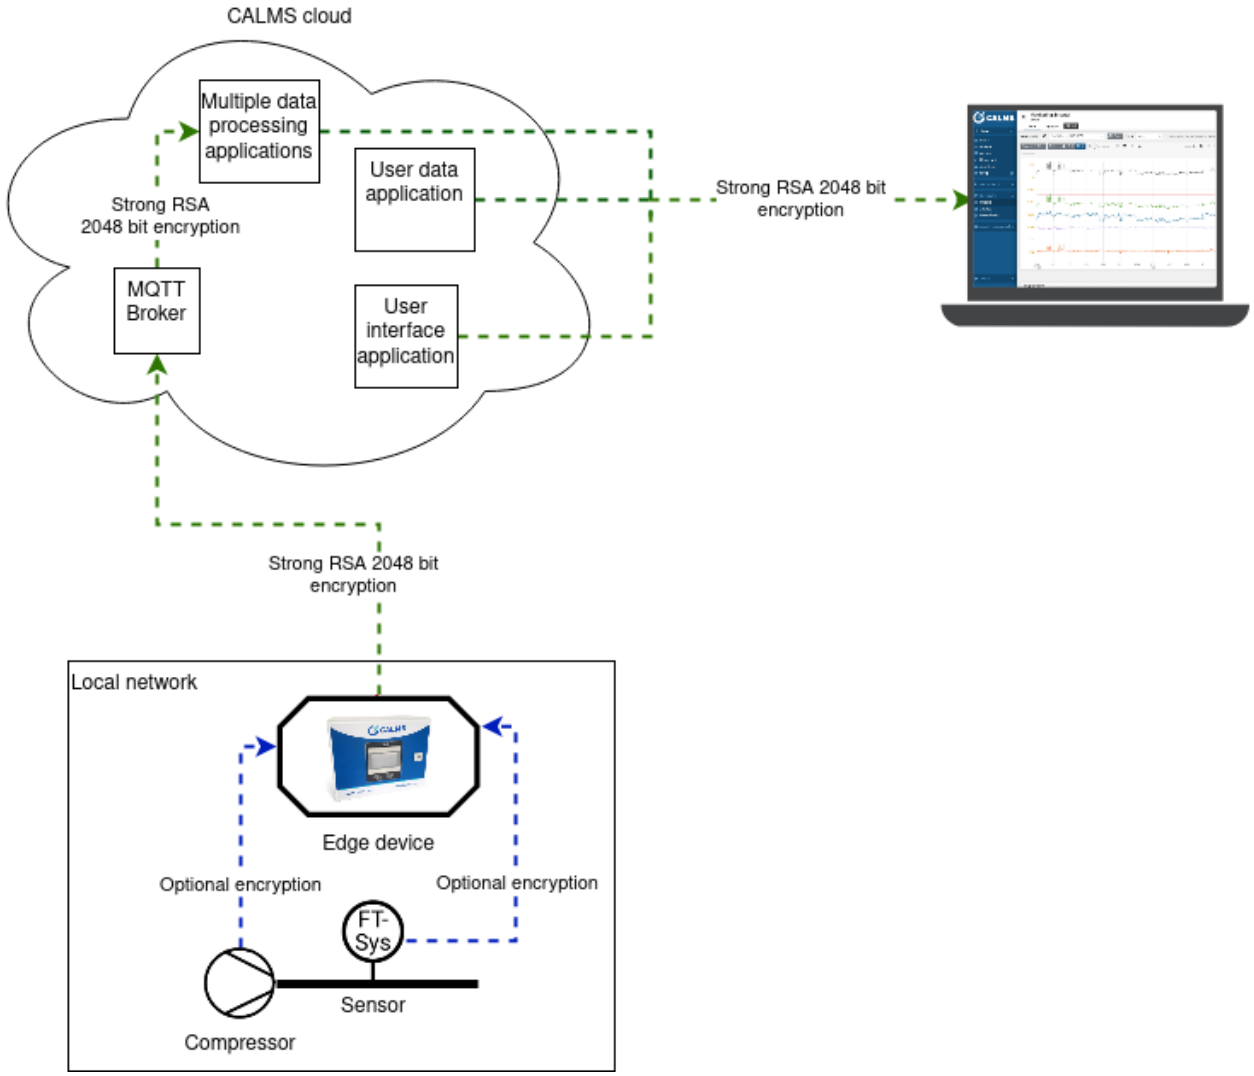

Data Logger

Data Logger to sensor communication

The data logger (also known as the Data Acquisition Module) is an IIOT device responsible for getting sensor values from customers measurement sites, encrypting them and sending to the cloud.

Within the on-site local network, the data logger collects data from sensors using encryption provided by sensor manufactures. In most cases, they do not provide any encryption and but this does not present a significant attack vector, because any attacker would need physical access to the data logger and/or sensors to intercept this data.

Access to data logger

After we capture the data, it is stored on the data logger. Access to this device is limited to a physical password protected terminal and SSH remote connection over internal VPN using authentication certificates. These certificates are issued only to a few of the vendors employees, who are responsible for device configuration, problem troubleshooting and regular software updates.

Data Logger to cloud communication

When data logger is connected to the internet (usually over GSM provider), it sends data to MQTT broker in the cloud using RSA 2048-bit encryption over MQTT protocol, all within the internal VPN. The data logger and MQTT broker authenticate each other using short-lived m-TLS certificates that are automatically regularly updated.

Data Acquisition System Cloud

Access to the cloud

The cloud infrastructure uses multiple mechanisms and layers to provide maximum security and prevent an attacker from accessing or modifying the data. These mechanisms have been developed and tested by many other cloud service providers, such as Google and Amazon. Access to the Data Acquisition System cloud infrastructure is severely limited, controlled and monitored.

Microservice architecture

The Data Acquisition System software is composed of multiple programs that can run independent of the others. This limits a potential attacker that has gained access to a part of the application only to that compromised part and not the whole application. For example, the Data Acquisition System isolates user data in a separate microservice, so it can be protected from a security incident in other parts of the application.

Isolation

All of the microservices (parts of the Data Acquisition System application) are orchestrated by Kubernetes, an advanced system for managing cloud infrastructure. Within Kubernetes, every part of

the application is isolated and enclosed in a separate Docker container. To an attacker, it would seem like every program runs on a separate computer and cannot access other parts of our application. To access another microservice is almost as hard as compromising another computer. To further improve security, the vendor made sure that those containers only have programs needed for their work. So if an attacker is able to gain access to one of the containers, this would be equivalent to gaining access to a computer without a display, keyboard, mouse or an operating system.

User application

Passwords and authentication

For user authentication the Data Acquisition System uses OAuth2 architecture based on a user chosen password and cryptographically signed JWT token. Before being stored, passwords are hashed using BCrypt HMAC algorithm, so in extreme case where an attacker gains access to the database, they will not be able to get retrieve passwords users may use in other applications.

Permissions

The application has an advanced permission system where each segment of the data is protected with a specific permission. Without granted correct permissions, a user cannot load any data. The system also logs every important user action so data does not get lost, deleted or modified without traces.

Preventing security leaks in software development

To ensure high quality and security of the product, the vendor has multiple practices in place, designed to identify security vulnerabilities before they are exploited. The vendor runs automated unit and end-to-end tests of their application, and reviews the code of developers alongside regular security reviews.

Backups

Backups are handled in multiple layers. Every data logger has its own internal backup of a few gigabytes. The age of the oldest data in the device backup depends on number of channels and sampling interval.

Another backup layer is a backup of the whole application and history of all the channels and devices. Such backups are made every 12 hours and together with data logger backups, the vendor can recover the whole system with no data loss in case of major data loss or other severe incident within the application.